

Coded Cooperative Data Exchange Problem for General Topologies

Mira Gonen

Department of Mathematics
Bar-Ilan University
gonenm1@math.biu.ac.il

Michael Langberg

Department of Mathematics and Computer Science
The Open University of Israel
mikel@open.ac.il

Abstract—We consider the *coded cooperative data exchange problem* for general graphs. In this problem, given a graph $G = (V, E)$ representing clients in a broadcast network, each of which initially hold a (not necessarily disjoint) set of information packets; one wishes to design a communication scheme in which eventually all clients will hold all the packets of the network. Communication is performed in rounds, where in each round a single client broadcasts a single (possibly encoded) information packet to its neighbors in G . The objective is to design a broadcast scheme that satisfies all clients with the minimum number of broadcast rounds.

The coded cooperative data exchange problem has seen significant research over the last few years; mostly when the graph G is the complete broadcast graph in which each client is adjacent to all other clients in the network, but also on general topologies, both in the fractional and integral setting. In this work we focus on the integral setting in general undirected topologies G . We tie the data exchange problem on G to certain well studied combinatorial properties of G and in such show that solving the problem exactly or even approximately within a multiplicative factor of $\log |V|$ is intractable (i.e., NP-Hard). We then turn to study efficient data exchange schemes yielding a number of communication rounds comparable to our intractability result. Our communication schemes do not involve encoding, and in such yield bounds on the *coding advantage* in the setting at hand.

I. INTRODUCTION

In this work we study the coded cooperative data exchange problem for general graphs. An instance to the problem consists of an undirected graph $G = (V, E)$ representing a communication network (in which each node of G represents a client, and edges in G represent client pairs that can communicate with each other), a parameter k representing the number of information packets $X = \{x_1, \dots, x_k\}$ to be transmitted over the network, and a set $\{X_i\}_{i \in V}$ of subsets of X representing the set of packets available at each client node $v_i \in V$ in the initial stage of the transition. The objective is to design a communication scheme in which, eventually, all nodes of the network will obtain all k packets. Loosely speaking, in each round of the communication scheme, a single node broadcasts a single (possibly encoded) packet to all its neighbors in G . The goal is to find a communication scheme in which the number of communication rounds is minimum.

The coded cooperative data exchange problem has seen significant research over the last few years. The problem was

introduced by El Rouayheb et al. in [15], where data exchange over a *complete* graph G was considered (in which each client can broadcast its messages to all other clients in G). In [15] certain upper and lower bounds on the optimal number of transmissions needed was established. In a subsequent work, Sprinston et al. [17] continue the study of complete graphs G and present a (randomized) algorithm that with high probability achieves the minimum number of transmissions, given that the packets are elements in a field F_q with q large enough. Ozgul et al. [13] study a variant of the data exchange problem in which each client has a distinct broadcast cost and one wishes to minimize the cost of the transmission scheme after which all clients have obtained all information packets. In [13], optimal randomized linear encoding schemes are given for the problem at hand.

Communication in which *fractional* packets can be transmitted is addressed in the works of Courtade et al. in [3] (for general topologies G) and Tajbakhsh et al. [18], [19] (for the complete topology). In the fractional setting, packets are assumed to be divisible into chunks so that a fraction of a packet may be transmitted at any (fractional) round of communication; as apposed to the integral setting in which information packets are indivisible. In [3], [18], [19] it is shown that the fractional setting of the data exchange problem reduces to that of multicast network coding and can be efficiently solved in an optimal manner via linear programming and the concept of linear network coding, see e.g. [1], [6], [7], [9], [10].

Most related to our work is the work of Courtade et al. in [4] which focus on general topologies G in the setting of indivisible packets (the integral setting). [4] continue the paradigm of [3] which characterizes the data exchange problem as a family of *cut inequalities*, and present certain communication schemes that yield approximate solutions for an asymptotic number of packets k . Roughly speaking, [4] analyze a certain communication scheme in which each client transmits at a certain *fixed rate* over time, and obtain nearly optimal rate allocations (within an additive approximation of ϵk for general graphs, and $|V|$ for regular graphs). An important aspect in the analysis in [4] is the assumption that the number of packets k tends to infinity. A detailed comparison of the results of [4] with ours appears below at the end of Section I-A.

Most recently, Milosavljevic et al. [12] present a comprehensive study of data exchange over the complete topology in

This work was supported in part by ISF grant 480/08 and the Open University of Israel's research fund (grant no. 46114). Work done in part while Mira Gonen was at the Open University of Israel.

which one wishes to broadcast the components of a (jointly distributed) discrete memoryless multiple source. Efficient optimal *rate* schemes are presented for a number of side information models.

A. Our contribution

In this work we study the coded cooperative data exchange problem on general topologies. We focus on the combinatorial *integral* setting in which one assumes that packets are indivisible. Namely, we assume that each packet is a value from a given alphabet Σ , and in each communication round a single element of Σ is broadcasted by a client to its neighbors in G . The study of the indivisible integral setting, rises naturally in communication schemes in which dividing information packets to several chunks leads to undesirable overhead in communication (via scheduling issues or rate loss due to header information). Our work addresses the design and analysis of *efficient* algorithms that (approximately) solve the problem at hand. Throughout our work, we assume that the number of packets k is polynomial in the size of the network $|V|$. In this context, an efficient algorithm is one which is polynomial in the network size.

We start by tying the data exchange problem in general topologies G to certain well studied combinatorial properties of G . Specifically, we consider the *Dominating Set* problem (e.g., [8]) and its variants (to be defined in Section II), and show that they are closely related to the data exchange problem. Namely, we show that (i) a solution to the Dominating Set problem (or its variants) yields a (not necessarily optimal) solution to the data exchange problem, and (ii) an optimal solution to the data exchange problem yields a nearly optimal solution to the Dominating Set problem(s). Roughly speaking, these connections (together with others) imply two initial results. Primarily, that it is NP-Hard to find a solution to the data exchange problem in which the number of communication rounds is within a multiplicative factor of $\Omega(\log |V|)$ from the optimal. Secondly, that a conceptually simple data exchange algorithm, that does not involve encoding, based on the Dominating Set problem yields a number of communication rounds which is within a multiplicative factor of $O(k \cdot \log |V|)$ from the optimal.

The gap between the upper and lower bounds above is k (the number of distinct packets in the network) which may be of significant size. Reducing this gap is the main focus of our work. Roughly speaking, in this work we reduce the gap of k by analyzing our algorithm based on the Dominating Set problem(s). Our algorithm does not involve coding and in such yields bounds on the *coding advantage* in the setting of data exchange. Our detailed results are given below, which at times are the best possible (assuming standard tractability assumptions).

The paper is structured as follows. In Section II, we present the model and notation used throughout this work, including the several variants of the Dominating Set problem used in our analysis. In Section III, we prove that it is NP-Hard to approximate the data exchange problem on general

topologies within a multiplicative factor of $\Omega(\log n)$ (for any k polynomial in n). Here, $n = |V|$.

In Section IV, we present our algorithm for data exchange based on the Dominating Set problem and its variants. The algorithm we present is conceptually very simple and does not involve coding. As mentioned above, a naive analysis of our algorithm yields an approximation ratio of $O(k \cdot \log n)$, and the majority of this section is devoted to proving that the algorithm actually performs better.

In Section IV-B, we show that our algorithm is the best possible (assuming standard tractability) and has an approximation ratio of $O(\log n)$ (matching the lower bound of Section III) on instances in which each packet is initially present at a single client in G . This implies a coding advantage of $O(\log n)$ in such cases.

In Section IV-C, we study data exchange instances in which the underlying graph is regular (each client has the same number of neighbors). We show that the approximation ratio in this case is again better than $O(k \cdot \log |V|)$ and depends on the average number \bar{d} of packets available at client nodes. Specifically we show that in this case the approximation ratio of our algorithm is $O\left(\frac{k}{k-\bar{d}}\right) \log n = O\left(1 + \frac{\bar{d}}{k-\bar{d}}\right) \log n$ (thus improving the factor of k in the naive analysis to $1 + \frac{\bar{d}}{k-\bar{d}}$). Notice, that for $\bar{d} = \Theta(k)$ (the case in which on average each client initially has a constant fraction of the packets) we obtain an approximation ratio that matches the bound of Section III. Our results imply a coding advantage of $O\left(1 + \frac{\bar{d}}{k-\bar{d}}\right) \log n$ in the cases at hand. Finally, in Section IV-D we study general graphs G with no restrictions and present an improved approximation ratio to that naively mentioned above.

We conclude our work by studying a refined version of our algorithm (still without encoding) in Section IV-E and by discussing future research directions in Section V.

Comparing our results with those in [4] is not straightforward. Courtade et al. [4] focus on the setting in which the number of packets k tends to infinity and may be significantly greater than the network size n . The setting of asymptotic k allows the design of algorithms which are efficient with respect to k but may be exponential in n . In our work we focus on the setting in which k is polynomially bounded by n , and obtain communication schemes that can be designed efficiently in time polynomial in the network size n . In addition, [4] focus on the case in which every client initially holds a constant fraction of the k information packets;¹ and in this setting study additive approximations. In this work, we study multiplicative approximations, and our assumptions (if any) on the packet distribution are of different nature.

II. MODEL DEFINITION AND PRELIMINARIES

A. Coded Cooperative Data Exchange Problem

We start by defining the Coded Cooperative Data Exchange Problem for General Graphs. Let $G = (V, E)$ be a given

¹The precise formulation in [4] is phrased in terms of “well behaved” packet distributions; i.e., the asymptotic (in k) empirical probability that a client (or set of clients) holds a certain number of packets.

undirected graph with $V = \{v_1, \dots, v_n\}$. Let $X = \{x_1, \dots, x_k\}$ be a set of packets to be delivered to the n clients belonging to the set V . The packets are elements of a finite alphabet which will be assumed to be a finite field F_q . At the beginning, each client v_i knows a subset of packets denoted by $X_i \subseteq X$, while the clients collectively know all packets in X . We denote by $\bar{X}_i = X \setminus X_i$ the set of packets required by client v_i . For each client (vertex in G) v_i let $d_{v_i} = |X_i|$ be the number of packets it holds, let $\bar{d} = \sum_{v \in V} d_v / n$ be the average number of packets present at vertices of G , and let $d = \max_{v \in V} d_v$ be the maximum number of packets that any client holds. We will use these parameters in our analysis.

Each client may transmit packets to its neighbors in G via a lossless broadcast channel capable of transmitting a single element in F_q . The data is transmitted in communication rounds, such that at round i one of the clients, say v , broadcasts an element $x \in F_q$ to all its neighbors in G . The transmitted information x may be one of the original packets in X_j , or some encoding of packets in X_j and the information previously transmitted to v_j .

Our goal is to devise a scheme that enables each client $v_i \in V$ to obtain all packets in \bar{X}_i (and thus in X) while minimizing the total number of broadcasts. This work focuses on the *integral* (i.e., *scalar*) setting in which each broadcast consists of a single element of F_q . We denote by NC the minimum number of (integral) broadcasts needed to satisfy the given instance to the coded cooperative data exchange problem at hand. In this work we connect the value of NC with other well studied combinatorial operators on G defined below.

Throughout our work, we assume that the number of packets k is polynomial in the size of the network $|V|$ (i.e., $k \leq |V|^c$ for some constant c). In this context, we say an algorithm is *efficient* if its running time is polynomial in the network size.

B. The Self Dominating Set problem

Given an undirected graph $G = (V, E)$, a *self dominating set* of G is a subset of vertices S such that every $v \in V$ is connected to some vertex $s \in S$ by an edge $(s, v) \in E$. In such a case we say that $v \in N(s)$ where $N(s) = \{v \mid (s, v) \in E\}$. The self dominating set problem is closely related to the standard dominating set problem, e.g. [8], on which we elaborate below. The minimum size of a self dominating set in G is denoted by DS^+ . A self dominating set S with a corresponding induced subgraph that is connected is referred to as a *connected self dominating set*. Denote by CDS^+ the size of a minimum connected self dominating set in G . We will show below that computing (or approximating) any of the values mentioned above (i.e., DS^+, CDS^+) is NP-Hard.

In this work we will also be interested in a fractional version of the Self Dominating Set problems expressed by the following linear program. Given a graph $G = (V, E)$, find a set of capacities $C = \{c_v \mid v \in V\}$ (where for each $v \in V$, c_v is the capacity of vertex v) such that $\sum_{v \in V} c_v$ is minimum, and $\forall v \in V$ it holds that $\sum_{u \in N(v)} c_u \geq 1$. The above is equivalent to the solution of the following LP:

$$\begin{aligned} &\text{Minimize} && \sum_{v \in V} c_v \\ &\text{subject to} && \sum_{u \in N(v)} c_u \geq 1, \forall v \in V \\ &&& 0 \leq c_v \leq 1, \forall v \in V \end{aligned}$$

Let DS_f^+ denote the minimum value of the linear program above. By considering integral values of c_v , it is straightforward to establish that $DS_f^+ \leq DS^+$.

As we will see, at times we would like to “cover” each vertex in G more than once by our self dominating sets S . We thus consider the integer and fractional k Self Dominating Set problems as well. Below we phrase the fractional version, with optimum denoted by $(k - DS^+)_f$, the integer variant is obtained by setting $c_v \in \{0, 1\}$ and its optimum will be denoted by $k - DS^+$:

$$\begin{aligned} &\text{Minimize} && \sum_{v \in V} c_v \\ &\text{subject to} && \sum_{u \in N(v)} c_u \geq k, \forall v \in V \\ &&& 0 \leq c_v \leq 1, \forall v \in V \end{aligned}$$

Finally, as we will see, to connect the cooperative data exchange problem with the notion of dominating sets in G , we will need to specify the “cover” requirement explicitly for each vertex v . We refer to this variant as the *Augmented- k -Fractional Self Dominating Set* problem. Here, we solve the same linear program with the exception that each vertex needs to be covered at least $k - d_v$ times (the use of the parameter d_v that was defined previously to be the number of initial packets present at v is not occasional).

$$\begin{aligned} &\text{Minimize} && \sum_{v \in V} c_v \\ &\text{subject to} && \sum_{u \in N(v)} c_u \geq k - d_v, \forall v \in V \\ &&& 0 \leq c_v \leq 1, \forall v \in V \end{aligned}$$

We denote by $A - (k - DS^+)_f$ the optimal solution to the linear program above. Note that the above is an augmented version of the k fractional self dominating set problem when there is an initial solution $\{d_v\}$ and we wish to *augment* it to a full solution by using values of $\{c_v\}$.

Some observations and related work expressing the relationships between the notions defined above are in place:

Lemma 1 $(k - DS^+)_f = k \cdot DS_f^+$.

Proof: Any solution $\{c_v\}$ to DS_f^+ implies a solution $\{c_v^k\} = \{k \cdot c_v\}$ to $(k - DS^+)_f$ and visa versa. ■

Note that the above lemma is not valid for the integral versions of the problems, namely $k - DS^+ \neq k \cdot DS^+$. E.g., it is not hard to verify that the 2 by 3 complete bipartite graph $(K_{2,3})$ with an additional edge between the two vertices in the 2-size side has $2 - DS^+ = 3$ and $DS^+ = 2$.

Lemma 2 Defining the parameters d_v to be equal to $|X_i|$ for every $v_i \in V$, it holds that $A - (k - DS^+)_f \leq NC$.

Proof: Consider any solution to the coded cooperative data exchange problem. For every vertex $v \in V$, let c_v be the number of times v transmitted information during the execution of the solution at hand. By our definitions $\sum c_v \geq NC$. We now show that $\{c_v\}$ is also a solution to $A - (k - DS^+)_f$. Namely, consider any $v \in V$ that is missing $k - d_v$ packets in our data exchange problem. It must be the case, that during

the process of communication it received at least $k - d_v$ broadcasts, as otherwise it could not be able to obtain all k packets after the communication process. Thus it holds that $\sum_{u \in N(v)} c_u \geq k - d_v$ as desired. ■

Lemma 3 *Let $\{d_v\}$ be the set of weights in the augmented k -dominating set problem, and let $d = \max_{v \in V} d_v$. Then*

$$(k - d) \cdot DS_f^+ \leq A - (k - DS^+)_f \leq NC.$$

Proof: The right inequality follows from Lemma 2. For the left inequality, we notice that each solution to the fractional augmented k self dominating set problem is a fractional solution to the $(k - d)$ self dominating set problem. Namely, let $\{c_v\}$ be the capacities of an optimal solution to the fractional augmented k self dominating set problem. Then for all v it holds that $\sum_{u \in N(v)} c_u \geq k - d_v \geq k - d$. Therefore, $\{c_v\}$ is a solution to the fractional $(k - d)$ self dominating set problem. Now using Lemma 1, we obtain:

$$(k - d) \cdot DS_f^+ = ((k - d) - DS^+)_f \leq A - (k - DS^+)_f.$$

■

C. The (standard) dominating set problem

We now address the standard dominating set problem, which slightly differs from the previously defined *self* dominating set problem. Given an undirected graph $G = (V, E)$, a (standard) dominating set of G is a subset of vertices S such that every $v \in V$ is either in S or connected to some vertex $s \in S$ by an edge $(s, v) \in E$. The minimum sized dominating set in G is denoted by DS . The fractional variant of the dominating set problem is expressed by the following linear problem:

$$\begin{aligned} &\text{Minimize} && \sum_{v \in V} c_v \\ &\text{subject to} && \sum_{u \in N(v) \cup \{v\}} c_u \geq 1, \forall v \in V \\ &&& 0 \leq c_v \leq 1, \forall v \in V \end{aligned}$$

We denote the optimal solution to the linear problem above by DS_f . Clearly, it holds that $DS \leq DS^+$ and that $DS_f \leq DS_f^+$.

As before, one can define the connected variant of the dominating set problem, and the k -dominating set problem. We denote the optional values in these cases as CDS for the connected variant, $k - DS$ for integral k -dominating set, and $(k - DS)_f$ for fractional k -dominating set. As in Lemma 1 we have that:

Lemma 4 $(k - DS)_f = k \cdot DS_f$.

The following lemma that constructively connects between dominating sets and their connected variant was proven in [5].

Lemma 5 ([5]) *Given any dominating set D , one can efficiently construct a connected dominating set D' with $|D'| \leq \frac{4}{3} \cdot |D|$. Specifically, for every connected graph $G = (V, E)$ it holds that $CDS \leq \frac{4}{3} \cdot DS$.*

It is NP-Hard to estimate the size of the minimum dominating set of a given graph G up to a multiplicative factor of $\Omega(\log |V|)$ [14]. Notice that if $CDS > 1$, then $CDS^+ = CDS$, (and in general $CDS^+ \leq CDS + 1$) so finding CDS ,

and CDS^+ (and also approximating them beyond a ratio of $\Omega(\log |V|)$) is also NP-hard. Lemma 5 and the definition of the self dominating set problem imply the following lemma which connects DS , DS^+ , CDS , and CDS^+ :

Lemma 6

$$\frac{4}{3}DS + 1 \geq CDS + 1 \geq CDS^+ \geq DS^+ \geq DS.$$

Lemma 6 implies that all the values DS , DS^+ , CDS , and CDS^+ are all approximately (up to constant factors) the same size.

III. INTRACTABILITY RESULTS

In this section we show that the coded cooperative data exchange problem is hard to approximate within a multiplicative factor of $c \log |V|$, for some $c > 0$, for every value of k . We use the fact that it is NP-hard to estimate DS within a multiplicative factor of $c \log |V|$, for some $c > 0$ [14]. We first show our hardness for $k = 1$. We then turn to the case of general k (polynomial in n).

Lemma 7 *The coded cooperative data exchange problem with $k = 1$ is NP-hard to approximate within $c \log |V|$, for some $c > 0$.*

Proof: We show that, essentially, the coded cooperative data exchange problem when $k = 1$ is equivalent to the connected dominating set problem. Namely, consider any (connected) instance $G = (V, E)$ of the dominating set problem and construct an instance to the data exchange problem which includes the network G and a single node $v_0 \in V$ that holds the (single) message x_1 . We show that the number of rounds in the optimal solution to the data exchange instance at hand NC is approximately the size of the minimum connected dominating set size CDS of G . Specifically

$$CDS \leq NC \leq CDS + 1$$

Consider an optimal solution to the data exchange problem. Notice that, as each edge has unit capacity, once there is only a single message x_1 to be broadcasted throughout the network, no encoding is needed. Thus, any solution to the data exchange problem will correspond to a series of broadcasts of message x_1 at certain nodes of the network. As there is only a single message, it also holds that no vertex needs to broadcast more than once. Let S be the set of vertices that performed a broadcast. The size of S is exactly the value of NC on the instance at hand. In addition, as every vertex $v \in V$ has received x_1 , it holds that either $v \in S$ or v is connected to S . This implies that S is a connected dominating set in G .

For the opposite direction, notice that any connected dominating set S in G implied a broadcast scheme for the data exchange problem. If v_0 is in S , then consider a broadcasting scheme that transmits according to a Breadth First Search (BSF) starting from v_0 in the subgraph induced by S . It is not hard to verify that such a scheme will use $|S|$ broadcasts and eventually will transmit x_1 to all the network. Namely, let (v_0, v_1, v_2, \dots) be a BSF ordering from v_0 on the vertices




Figure 1. Illustration of the graph G' of Lemma 8.

of S . The message $x_1 \in X$ can be transmitted from v_0 to all nodes in V using the ordering $(v_0, v_1, v_2 \dots)$. Specifically, our ordering implies that node v_j holds the message x_1 after nodes $\{v_0, v_1, \dots, v_{j-1}\}$ transmit and, as S is dominating, all nodes will eventually receive the message x_1 . If v_0 is not in S then it is connected to $s \in S$, so we can add v_0 to S and still have a connected dominating set (and now use the scheme described in the last paragraph). All in all, the resulting scheme will have $CDS + 1$ broadcasts.

As it is NP-hard to approximate CDS within a multiplicative factor of $c \log n$ for some universal constant $c > 0$ on graphs of size n for which CDS depends on n (this follows directly from [14] and Lemma 6); it holds that the same is true for the parameter NC under study. ■

Note that the proof of Lemma 7 is also valid for $k = 1$ in the specific case when only one vertex holds the information. This implies that our upper bound for the case of disjoint sets of messages discussed in Section IV-B is tight.

We now show that our hardness result holds for every k by (again) presenting a reduction from the dominating set problem. Given an instance $G = (V, E)$ to the dominating set problem, we construct the following graph $G' = (V', E')$ for the coded cooperative data exchange problem. G' has k copies of G , and a new vertex v , such that v is connected to a vertex u_i in each copy G_i of G . Figure 1 illustrates G' . All vertices u_i know all messages, v knows no message, and for each G_i all vertices in G_i besides u_i know all messages besides the i 'th one.

Lemma 8 *If $DS(G) \leq \alpha$ then $NC(G') \leq k \cdot \frac{4}{3} \cdot (\alpha + 1)$, and if $DS(G) > \beta$ then $NC(G') > k \cdot \beta$.*

Proof: Assume that $DS(G) \leq \alpha$. Then the following is a transmission algorithm in $k \cdot \frac{4}{3} \cdot (\alpha + 1)$ communication rounds. Let U_i be a minimum connected dominating set of G_i . For all $1 \leq i \leq k$, as only one message needs to be broadcasted throughout G_i , one may design a broadcast scheme to satisfy all nodes in G_i based on the connected dominating set U_i exactly as in the proof of Lemma 7. It holds (via Lemma 6)

that

$$\begin{aligned} NC(G') &\leq \sum_{i=1}^k (|U_i| + 1) \leq \sum_{i=1}^k \frac{4}{3} \cdot DS(G_i) + 1 \\ &= k \cdot \left(\frac{4}{3} \cdot DS(G) + 1 \right) \leq k \cdot \frac{4}{3} \cdot (\alpha + 1). \end{aligned}$$

Now assume that $NC(G') \leq k \cdot \beta$. We first show that $CDS(G_i) \leq NC(G_i)$. This follows by the proof of Lemma 7, as any communication scheme in G_i only needs to communicate a single message x_i from u_i to the vertices of G_i (recall that all vertices in G_i know all the messages in $X \setminus \{x_i\}$). Now, it also holds that $DS(G_i) \leq CDS(G_i)$ and that $NC(G_i) \leq NC(G')/k$, thus we obtain $DS(G) = DS(G_i) \leq NC(G_i) \leq \beta$. All in all, we now conclude that estimating $NC(G')$ within a multiplicative factor of $O(\log n)$ will imply such an estimate for $DS(G)$. ■

Lemma 7, Lemma 8 and the hardness of computing DS specified in [14] imply the following theorem:

Theorem 1 *The coded cooperative data exchange problem is NP-hard to approximate within $c \log |V|$, for some $c > 0$, for every value of k polynomial in $|V|$.*

IV. APPROXIMATION ALGORITHM

In this section we give an approximation algorithm for the coded data exchange problem and analyze its approximation ratio. In the first subsection we present the approximation algorithm. In the second subsection we analyze the quality of the algorithm on a number of graph families or initial packet allocations, and show that for these instances the approximation ratio of the given algorithm matches (or comes close to matching) the results given in the previous section. In the third subsection we extend our analysis to the general case.

A. The Algorithm

The following lemma introduces an approximation algorithm for the cooperative data exchange problem.

Lemma 9 *Given a connected dominating set D of G one can efficiently solve the cooperative data exchange problem in $k \cdot (|D| + 1)$ communication rounds. Specifically, $NC \leq k \cdot (CDS + 1)$.*

Proof: The proof follows that given in Lemma 7. Let D be a connected dominating set in G . Let s_i be an arbitrary node holding message x_i . Assume that s_i is a node in D . Let $(s_i, v_1, v_2 \dots)$ be a BSF (Breadth First Search) ordering from s_i on the vertices of D . The message $x_i \in X$ can be transmitted from s_i to all nodes in V using the ordering $(s_i, v_1, v_2 \dots)$. Specifically, our ordering implies that node v_j holds the message x_i after nodes $\{s_i, v_1, \dots, v_{j-1}\}$ transmit and, as D is dominating, all nodes will eventually receive the message x_i . All in all, transmission of the k messages will take $k \cdot CDS$ communication rounds. If s_i is not in D , then an additional round of communication is required for each message in order for it to reach the set D . ■

Since the problem of finding a minimum connected dominating set is NP-hard, we need to show how to approximately find such a set (efficiently). Roughly speaking, we will find a connected dominating set in our network G by first solving the fractional dominating set problem, by then *rounding* the fractional solution to an integral one to obtain a standard dominating set of G (see e.g., [2], [8], [11], [16]), and by finally modifying the dominating set to a connected one via Lemma 5. All in all, this (well studied) scheme will yield a connected dominating set D of size at most $c \log n \cdot DS_f$ for some universal constant $c > 0$.

Repeating the above more formally, given an instance G to the cooperative data exchange problem on general topologies, one can efficiently perform the following algorithm:

- 1) Solve the fractional dominating set problem on G to obtain a fractional solution $\{c_v^f\}$.
- 2) Change the fractional solution to an integral one $\{c_v\}$ corresponding to a dominating set D (via, e.g., [2], [8], [11], [16]).
- 3) Using D , construct a connected dominating set D' (Lemma 5) with $|D'| = O(|D|)$.
- 4) Broadcast the k source messages according to the procedure specified in Lemma 9 in $O(k|D'|) \leq O(k \log n \cdot DS_f)$ communication rounds.

The procedure above will yield a communication scheme with at most $O(k \log n \cdot DS_f)$ communication rounds. To understand the quality of the algorithm, one must express the size NC (or at least bound it from *below*) by an expression which can be easily compared with the bound $O(k \log n \cdot DS_f)$. For example, consider an instance to the data exchange problem in which $d = \max_{v \in V} d_v < k$ (here, for all $v_i \in V$ $d_{v_i} = |X_i|$). We have seen via Lemma 3 that $NC \geq (k - d) \cdot DS_f^+ \geq (k - d) \cdot DS_f \geq DS_f$. Thus, on these instances we obtain a solution to the data exchange problem that is within a multiplicative factor of $O(k \log n)$ from the optimal solution. It is also not hard to see (we do this implicitly in Section IV-E) that even if $d = \max_{v \in V} d_v = k$ a slight variant to our algorithm yields a solution which is within a multiplicative factor of $O(k \log n)$ from the optimal solution. The next sections attempt to improve this ratio to better match the hardness results presented in Section III. Specifically, we show that the factor of k in the ratio $O(k \log n)$ can be reduced or in cases removed.

B. Disjoint Sets of Messages

In this subsection we analyze our approximation algorithm for the case that for each two nodes v, u it holds that $X_v \cap X_u = \emptyset$. Note that this includes the case where only one node holds all the information, and all other nodes have no information. Namely, for some $v \in V$, $X_v = X$, and for all $u \neq v$ $X_u = \emptyset$. For this case we are able to improve over the lower bound presented in Lemma 3.

Lemma 10 $NC \geq k \cdot DS_f$.

Proof: We show that a solution to the Coded Cooperative Data Exchange problem induces a solution to the k Dominating Set

problem. For every vertex $v \in V$ define c_v to be the number of packets transmitted by v during an optimal data exchange protocol. It holds that for every vertex v the sum of capacities c_u of all $u \in N(v) \cup \{v\}$ is at least k . This is true since each node v must send at least $|X_v|$ packets (as no other node holds the packets in X_v and they must eventually reach the entire network), and receive at least $k - |X_v| = |\bar{X}_v|$ packets. Therefore $\sum_{u \in N(v) \cup \{v\}} c_u = \sum_{u \in N(v)} c_u + c_v \geq |\bar{X}_v| + |X_v| = k$. Thus $(k - DS)_f \leq k - DS \leq NC$. Finally, by Lemma 4 it holds that $k \cdot DS_f = (k - DS)_f$. ■

As our algorithm gives a communication scheme with at most $O(k \log n \cdot DS_f)$ rounds we conclude:

Theorem 2 *If for every two nodes v, u it holds that $X_v \cap X_u = \emptyset$, the cooperative data exchange problem on general topologies can be efficiently solved within an approximation ratio of $O(\log n)$. Moreover, in such cases it holds that*

$$k \cdot DS_f \leq NC \leq k \cdot \left(\frac{4}{3} \cdot DS + 1 \right) \leq O(k \log n) DS_f.$$

As our algorithm does not involve coding, this implies that the coding advantage is $O(\log n)$.

C. Regular Graphs

In this subsection we show that if the given graph is regular our approximation algorithm has a $(1 + \bar{d}/(k - \bar{d})) \cdot O(\log n)$ approximation ratio. As before, we start by giving a lower bound for NC in this case. Let G be a Δ regular graph, and let $\bar{d} = \frac{1}{n} \sum d_v$, then it holds that

Lemma 11 $(k - \bar{d}) DS_f \leq NC$.

Proof: Consider the optimal communication scheme for the data exchange problem. Since every vertex v must receive at least $k - d_v$ messages, the total number of *edge* transmissions over the network is at least $\sum_{v \in V} k - d_v$ (here we are counting a single broadcast over r edges as r “edge”-transmissions). Since each broadcast may transmit over at most Δ vertices it follows that

$$NC \geq \frac{\sum_{v \in V} k - d_v}{\Delta} = \frac{n(k - \bar{d})}{\Delta} \geq (k - \bar{d}) DS_f^+ \geq (k - \bar{d}) DS_f.$$

For the second inequality, notice that one can obtain a fractional self dominating set by setting $c_v = \frac{1}{\Delta}$ for each $v \in V$. This implies that $DS_f^+ \leq \frac{n}{\Delta}$. The last inequality holds by definition of DS_f and DS_f^+ . ■

The following theorem follows from the above lemma:

Theorem 3 *The cooperative data exchange problem on regular topologies has a $(1 + \bar{d}/(k - \bar{d})) \cdot O(\log n)$ approximation ratio. Specifically,*

$$(k - \bar{d}) \cdot DS_f \leq NC \leq k \cdot \left(\frac{4}{3} \cdot DS + 1 \right) \leq O(k \log n) DS_f.$$

As our algorithm does not involve coding, this implies that the coding advantage is $O\left(\left(1 + \frac{\bar{d}}{k - \bar{d}}\right) \log n\right)$.

Proof: By Lemma 11 we have that

$$(k - \bar{d}) \cdot DS_f \leq NC.$$




Figure 2. Illustration of the graph G for our counter example to Lemma 11 on general (non-regular) graphs.

All in all, we obtain a solution with cost

$$\begin{aligned}
 O(\log n) \cdot k \cdot DS_f &= O(\log n) \cdot \frac{k}{k - \bar{d}} \cdot (k - \bar{d}) DS_f \\
 &\leq O(\log n) \cdot NC \cdot \frac{k}{k - \bar{d}} \\
 &= O(\log n) \cdot NC \cdot (1 + \bar{d}/(k - \bar{d})).
 \end{aligned}$$

■

D. General Case

In this subsection we analyze the quality of our approximation algorithm for any instance G . We first give an example that shows that our lower bound for regular graphs stated in Lemma 11 of $(k - \bar{d}) \cdot DS_f$ does not hold for general graphs.

1) *Example, complementing Lemma 11:* We present a (general, non-regular) graph G for which the lower bound stated in Lemma 11 of $(k - \bar{d}) \cdot DS_f$ does not hold (even in an approximate manner). Consider a graph G that consists of two parts: The first part is a set of m (disjoint) cliques of size k . In each clique, for each message i (between 1 to k) there is exactly 1 vertex missing message i and having all the rest. The second part consists of a clique of size mk in which one vertex has all the information, and all the rest do not have any message. Figure 2 illustrates G . The value of an optimal scheme for data broadcast on the first part of G is $2m$ since for each clique two messages must be sent (One client broadcasts an arbitrary message. This will cause another client to have all of X , and it broadcasts the sum of all the messages in X over F_q). The value of an optimal scheme for data broadcast on the second part is obviously k (just perform k broadcasts from the single node that has all of X). Thus NC is $2m + k$. Now, it is not hard to verify that DS_f of the first part is m (1 for each clique), and DS_f of the second part is 1. Therefore DS_f is $m + 1$. Moreover, $\bar{d} = \frac{(k-1)km + mk}{2km} = \frac{k}{2}$, so $(k - \bar{d})DS_f = \frac{k}{2}(m + 1)$. Therefore, for large $m \gg k$ we

get that $NC \sim 2m$, and $(k - \bar{d})DS_f \sim \frac{km}{2}$, so we get a gap of approximately $k/4$ w.r.t. the assertion of Lemma 11.

2) *Generalizing Lemma 11:* We use Δ to denote the maximum degree of G and δ to denote the minimum degree of G . We generalize Lemma 11 to the case of general graphs:

Lemma 12 $\frac{\delta}{\Delta}(k - \bar{d})DS_f \leq NC$.

Proof: As in the proof of Lemma 11, the total number of edge transmissions over the network is at least $\sum_{v \in V} k - d_v$. Since each message can be transmitted to at most Δ vertices it follows that

$$\begin{aligned}
 NC &\geq \frac{\sum_{v \in V} k - d_v}{\Delta} = \frac{n(k - \bar{d})}{\Delta} \\
 &\geq \frac{\delta}{\Delta}(k - \bar{d})DS_f^+ \geq \frac{\delta}{\Delta}(k - \bar{d})DS_f. \quad (1)
 \end{aligned}$$

In the setting at hand, the second inequality is valid since $\frac{n}{\delta}$ is an upper bound for DS_f^+ (i.e., one may set every c_v to be equal to $\frac{1}{\delta}$ to get a valid solution to the linear program defining DS_f^+). ■

We now conclude (recall that $d = \max_{v \in V} d_v$):

Theorem 4 *The cooperative data exchange problem on general topologies has an approximation ration and coding advantage of*

$$O(\log n) \cdot \min \left\{ \left(1 + \frac{d}{k - \bar{d}} \right), \frac{\Delta}{\delta} \left(1 + \frac{\bar{d}}{k - \bar{d}} \right) \right\}.$$

Proof: Using Lemma 3 we have that

$$(k - d) \cdot DS_f \leq (k - d) \cdot DS_f^+ \leq NC.$$

In addition, by Lemma 12 we have that

$$\frac{\delta}{\Delta}(k - \bar{d}) \cdot DS_f \leq NC, \quad (2)$$

Thus, the cost $O(\log n) \cdot k \cdot DS_f$ of our solution is at most:

$$O(\log n) \cdot NC \cdot \min \left\{ (1 + d/(k - d)), \frac{\Delta}{\delta} \cdot (1 + \bar{d}/(k - \bar{d})) \right\}.$$

■

E. A Tighter Upper Bound

We now present a refined version of our algorithm from Section IV-A. The algorithm we present will not yield improved asymptotic (in n) approximation ratios, however it yields improved communication schemes that at times may match those returned by the algorithm of Section IV-A and at times may be significantly better (depending on the instance at hand).

Roughly speaking, we improve the previous algorithm by taking into account the simple fact that it suffices to send each packet $x_i \in X$ only to those clients that do not hold it. Therefore we do not actually need to find a connected dominating set. Instead, we can do the following. Let V_i be the set of vertices holding information packet x_i . Let $\bar{V}_i = V \setminus V_i$. A minimum sized \bar{V}_i -self dominating set is a minimum sized set of vertices $S \subset V$ such that each vertex in \bar{V}_i has a neighbor




Figure 3. Illustration of G'_i to be used in the construction of DS_i .

in S , and each connected component of S intersects V_i . Using \bar{V}_i -self dominating sets we can refine our algorithm.

Assume first that we know how to find a minimum sized \bar{V}_i -self dominating set for each message x_i . Let $\{C_j^i\}_{j=1}^\ell$ be the set of connected components of such a minimum sized \bar{V}_i -self dominating set. Let w_j be an arbitrary vertex in $C_j^i \cap V_i$. To communicate x_i throughout C_j^i we use the following natural procedure: w_j sends x_i , and then each vertex in C_j^i that received x_i sends x_i . It immediately follows, after performing this process for each connected component C_j^i , that all vertices in G hold x_i . Moreover, the number of communication rounds used in this scheme is equal to the size of the \bar{V}_i -self dominating set. Let $DS_i = DS_i(G)$ denote the minimum size of a \bar{V}_i -self dominating set in G .

We now turn to approximating DS_i . We define the following graph $G'_i = (V'_i, E'_i)$ corresponding to our definition of a \bar{V}_i -self dominating set: $V'_i = V_i$, and $E'_i = E \cup (V_i \times V_i)$. Figure 3 illustrates the construction of G'_i .

Lemma 13 $CDS(G'_i) \leq DS_i(G) \leq CDS(G'_i) + 1$.

Proof: Let D be a minimum sized \bar{V}_i -self dominating set in G . Then by definition of DS_i it follows that D is a connected dominating set in G'_i , since every connected component of D in G has a vertex in V_i , and all vertices in V_i are connected in G'_i . Similarly, let \bar{D} be a minimum connected dominating set in G'_i . If \bar{D} includes a vertex in V_i then it follows that \bar{D} is also a \bar{V}_i -self dominating set in G . Otherwise (as we assume w.l.o.g. that G is connected) we add one vertex v in V_i to \bar{D} . Here, we take v to be any vertex in V_i , as they are all connected to \bar{D} . This completes the proof. ■

By Lemma 13 we can efficiently perform the following algorithm:

- 1) For $1 \leq i \leq k$ do:
 - a) Construct G'_i .
 - b) Using the algorithm specified in Section IV construct a connected dominating set D_i in G'_i via the corresponding fractional solutions, with $|D_i| = O(\log |V'_i| \cdot DS_f(G'_i))$.

- c) For each connected component of D_i in G broadcast x_i (according to the procedure specified in the discussion above) in $O(|D_i|)$ communication rounds.

All in all, the refined algorithm efficiently solves the data exchange problem in $O(\sum_i \log |V'_i| \cdot DS_f(G'_i))$ rounds of communication which is at most the number $O(k \log n \cdot DS_f(G))$ of rounds from the original algorithm. This follows since G is subgraph (in edges) of G'_i , and thus by definition $DS_f(G'_i) \leq DS_f(G)$. Thus, our refined algorithm is at least as good as that of Section IV-A and improves over it in cases in which DS_i is significantly smaller than the dominating set in G .

V. CONCLUDING REMARKS

In this paper, we consider the cooperative data exchange problem for general topologies G in the combinatorial *integral* setting. We establish both upper and lower bounds on the multiplicative approximation ratio that one may obtain efficiently by tying our problem to certain well studied combinatorial properties of G . Our achievability results are based on communication schemes that do not involve coding and in such imply bounds on the coding advantage of the problem at hand. Our results address the setting of undirected networks. Extending our results to the case of directed graphs (by studying directed analogs to dominating sets) involves modifications in our analysis and is subject to future research.

REFERENCES

- [1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung. Network Information Flow. *IEEE Transactions on Information Theory*, 46(4):1204–1216, 2000.
- [2] V. Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4:233–235, 1979.
- [3] T. Courtade, B. Xie, and R. Wesel. Optimal exchange of packets for universal recovery in broadcast networks. In *Proceedings of Military Communications Conference*.
- [4] T. A. Courtade and R. D. Wesel. On the minimum number of transmissions required for universal recovery in broadcast networks. In *Proceedings of Forty-Eighth Annual Allerton Conference on Communication, Control, and Computing*.
- [5] S. Guha and S. Khuller. Approximation algorithms for connected dominating sets. *Algorithmica*, 20:374–387, 1998.
- [6] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong. A Random Linear Network Coding Approach to Multicast. *IEEE Transactions on Information Theory*, 52(10):4413–4430, 2006.
- [7] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen. Polynomial Time Algorithms for Multicast Network Code Construction. *IEEE Transactions on Information Theory*, 51(6):1973–1982, June 2005.
- [8] D. S. Johnson. Approximation algorithms for combinatorial problems. *J. Comput. System Sci.*, 9:256–278, 1974.
- [9] R. Koetter and M. Medard. An Algebraic Approach to Network Coding. *IEEE/ACM Transactions on Networking*, 11(5):782 – 795, 2003.
- [10] S.-Y. R. Li, R. W. Yeung, and N. Cai. Linear Network Coding. *IEEE Transactions on Information Theory*, 49(2):371 – 381, 2003.
- [11] L. Lovasz. On the ratio of optimal integral and fractional covers. *SIAM Journal on Discrete Math*, 13:383–390, 1975.
- [12] N. Milosavljevic, S. Pawar, S. E. Rouayheb, M. Gastpar, and K. Ramchandran. Optimal deterministic polynomial-time data exchange for omniscience. *Manuscript; available on arxiv.org*.
- [13] D. Ozgul and A. Sprintson. An algorithm for cooperative data exchange with cost criterion. In *Proceedings of Information Theory and Applications Workshop (ITA)*.

- [14] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and sub-constant error-probability PCP characterization of NP. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing (STOC)*, pages 475–484, 1997.
- [15] S. El Rouayheb, A. Sprintson, and P. Sadeghi. On coding for cooperative data exchange. In *Proceedings of ITW*.
- [16] P. Slavyk. A tight analysis of the greedy algorithm for set cover. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing (STOC)*, pages 435–441, 1996.
- [17] A. Sprintson, P. Sadeghi, G. Booker, and S. El Rouayheb. A randomized algorithm and performance bounds for coded cooperative data exchange. In *Proceedings of ISIT*, pages 1888–1892.
- [18] S. Tajbakhsh, P. Sadeghi, and R. Shams. A generalized model for cost and fairness analysis in coded cooperative data exchange. In *Proceedings of International Symposium on Network Coding (NetCod)*.
- [19] S. Tajbakhsh, P. Sadeghi, and R. Shams. A model for packet splitting and fairness analysis in network coded cooperative data exchange.